

Spacecraft Autonomy and the Missions of Exploration

Richard J. Doyle
Jet Propulsion Laboratory
California Institute of Technology

Lead article submitted to *IEEE Intelligent Systems* for the
Special Issue on Autonomous Space Vehicles
R. Doyle, Guest Editor

Abstract

Practitioners of **artificial intelligence** are being engaged to help develop the next generation of **flight software** for **NASA missions**, in partnership with other computer scientists, mission designers, operations personnel, spacecraft engineers, systems engineers, software engineers, and scientists. The goal is **spacecraft autonomy**, an onboard system-level capability to make mission-relevant decisions about which actions are needed and which data is important, without the benefit of ongoing **ground support**. Success in developing more autonomous spacecraft is a key ingredient in the NASA vision to achieve the next phase of **space exploration**, characterized by many more space platforms operating at once, more effective use of limited communications resources, and bolder mission concepts involving direct *in situ* investigation of remote environments.

Introduction

The last three years have been ones of challenge and to a certain extent, vindication, for AI practitioners at NASA. These years have brought the opportunities that most of us always had in mind when we chose careers in the space program: the chance to contribute directly to the spacecraft missions NASA conducts, by deploying AI software not only for ground support but also directly on the space platform, software which would play an integral part in the concept and success of the missions.

The changes which enabled the emergence of these opportunities have their roots in the well-known “faster, better, cheaper” challenge issued within NASA by its Administrator, Daniel Goldin. Mission and spacecraft designers, flight project managers and technologists all have been asked to make thoughtful contributions towards new kinds of missions which utilize new technologies and manage risks in new ways. But the goal is not only to find ways to shorten mission development lifecycles and reduce launch and operations costs (the “faster, cheaper” parts), but also to initiate a new era of exploration characterized by sustained in-depth scientific studies at increasingly remote environments (the “better” part). Spacecraft autonomy has a specific and essential role to play in this view of NASA’s future mission set: the closing of planning, decision and control loops onboard the space platforms rather than through human operators on the ground, to not simply enhance, but to enable bolder and unprecedented space mission concepts.

An early achievement within NASA’s “faster, better, cheaper” paradigm was the recent Mars Pathfinder mission, with its endearing rover Sojourner. The method of landing at Mars was clearly new, and aggressive: more or less throwing the lander and rover at the planet within a cushion of airbags to absorb the impact. The technique proved an unqualified success, and it was only a matter of hours before the first images of the Martian surface were available on the Web, and soon thereafter Sojourner had crawled down a

deployment ramp to begin months of valuable scientific studies directly on the surface of Mars.

The rover employed some simple autonomy capabilities: Sojourner was able to terminate traverse activities by detecting an expected landmark, typically a rock, at the end of the traverse. This is a simple form of landmark-based navigation which will become increasingly critical for the much longer traverses planned for future Mars rover missions, where techniques based on dead reckoning will not scale. The rover utilized a laser-based system for detecting obstacles, painting a known pattern of laser light on nearby objects and interpreting the size and distortion of the pattern to infer the proximity and crude shape of obstacles. Sojourner's locomotion system is a fine example of achieving a kind of autonomy through engineering design. The system is extremely robust, allowing the rover to safely negotiate objects up to one-half of its own height, thereby rendering them non-obstacles and eliminating the need to actively characterize them and reason about how to avoid them.

The limited autonomy capabilities of the rover Sojourner are based on research and development work carried out several years ago at NASA, when the relevance of AI techniques for the missions was much less generally accepted. The landscape is different now. There may be still disagreement about what forms of autonomy are needed, and how best to go about developing and deploying these new capabilities, or even about what autonomy *is*, exactly, as NASA embraces a shift which is as much cultural as it is technological, but the importance of autonomy -- and the AI which underlies it -- for many of the future missions is readily apparent and agreed upon [Doyle 97, Muscettola et al 98].

This special issue on Autonomous Space Vehicles reports on much of the current work at NASA aimed at designing, developing, deploying and evaluating autonomy capabilities for space platforms. This lead article has the purpose of placing this exciting AI work fully in its NASA context, and specifically in the context of the future planned missions of exploration -- fascinating in their own right -- which require, in some cases, cry out, for autonomy.

The Strategic Value of Autonomy

Autonomy on space vehicles will have three forms of payoff for NASA: the reduction of mission costs (an example of "cheaper"), the more efficient use of always limited communications links between the ground and the space platform (an example of "faster"), and the enabling of whole new mission concepts, each involving some new form of loop-closing¹ onboard the remote vehicle (an example of "better").

AI, and in particular, model-based techniques have the potential to make cost reduction impacts across the entire NASA mission lifecycle: to allow constraints to be understood explicitly and quantitatively in the earliest mission concept design studies, to provide modeling languages and tools to capture appropriate knowledge in the first stages of detailed design, to be carried forward to the rest of the mission lifecycle, to contribute to new software engineering concepts and techniques for generating, testing and reusing autonomy software, and finally, and perhaps most obviously, to impact mission operations. The degree of success in reducing operations costs by migrating traditionally ground-based functions to the spacecraft, providing a more direct link between mission

¹ Examples are the control loop involved in landing on a small body like an asteroid, where the gravitational field is difficult to model, or the science planning loop between detection of a scientifically interesting and transient phenomenon and timely planning of focused observations to capture the phenomenon.

scientists and the space platform, and in general decoupling the space vehicle from the traditional form of ground support will most likely be the first criteria against which autonomy capabilities are evaluated. Certainly there must be a shift from a paradigm of large dedicated ground teams for each mission to smaller ground teams shared among several missions.

While the imperative of reducing mission lifecycle costs is easily understood, the greater strategic value of autonomy may in fact be elsewhere. For a long time, data collection technologies as embodied in sensors and instruments have been easily outstripping the capacity of data analysis techniques and technologies. The normal science data processing and analysis lifecycle for a NASA mission involves downlinking all raw data and assembling a ground-based archive, on which the mission science team and later the science community at large perform offline analysis, typically for years. With concomitant advances in data mining, image analysis and machine learning technologies on the one hand, and onboard computing technologies on the other, there is now the very real possibility of performing some forms of science data analysis onboard the spacecraft, in near real-time. Two advantages emerge from such an approach: transient opportunities which require the quick and reliable recognition of scientifically interesting events are captured (such opportunities clearly are lost in the normal course of delayed offline analysis), and more efficient and flexible use of the precious downlink resource is enabled, through downlink prioritization, and in some cases, the onboard construction of more compact, perhaps more useful science products from the raw data.

But perhaps the most exciting -- and important -- use for spacecraft autonomy is in the enabling of new kinds of missions, ones not previously within reach because they require the space platform to operate in an unprecedented closed-loop fashion in its environment. The greatest strategic payoff for autonomy is here, because the potential is nothing less than the launching of the next major phase of space exploration, beyond the reconnaissance missions which have already been completed (with great success). These future missions are to be characterized by sustained *in situ* scientific studies, with themes as compelling as the search for life in the universe.

Future Mars rover missions provide a good example of the need to close loops between science-related detection and mission planning. During long traverses from one pre-selected science site to another, the rover should be able to detect potentially significant scientific phenomena and halt the traverse, conducting preliminary analyses and waiting for further instructions.

Another new form of loop-closing involves constellation missions comprised of multiple space platforms. Here loop-closing takes the form of coordination among the platforms, which is most interesting when they carry different assets. An example from Earth orbit is the spaceborne detection of environmental hazards like forest fires or volcanic eruptions. The first satellite to detect such an event may not have the most appropriate instrument for studying it, but when it sends out an alert across an entire Earth-observing fleet, other instruments can be brought to bear, each platform making its own decision on whether and how to contribute to the study of the event.

These are just a few examples of future mission concepts where the contributions of autonomy will be as essential as those coming from any traditional form of spacecraft engineering or mission design expertise.

Components of Spacecraft Autonomy

The capabilities which contribute to spacecraft autonomy may be divided into six categories. These include (in no particular order), automated guidance, navigation and control, mission planning, scheduling and resource management, intelligent execution, model-based fault management, onboard science data analysis, and autonomy architectures and software engineering. Nearly all of these areas will be treated in depth in the articles which comprise this special issue.

- Automated guidance, navigation and control is the form of autonomy with the longest history and is what most spacecraft and mission people first (sometimes only) think of when asked about autonomy. The area includes target body characterization and orbit determination, maneuver planning and execution, precise pointing of instruments, landmark recognition and hazard detection during landing, and formation flying.
- Mission planning, scheduling and resource management addresses spacecraft activity planning from high-level mission goals, and activities replanned when science or engineering events occur. Planned activities are automatically checked against available spacecraft resources and hard temporal constraints from the mission timeline².
- Intelligent execution is about task-level execution, monitoring and control, contingency management, and overall coordination of spacecraft activities. The capability also provides a measure of protection against software failures.
- Model-based fault management comprises anomaly detection, fault diagnosis, and fault recovery. Through the use of model-based reasoning, reliable fault protection can be achieved without comprehensive space platform safing, loss of mission context³, or immediate ground intervention required when faults occur.
- Onboard science data processing includes trainable object recognizers and knowledge discovery methods applied to, among other objectives, prioritizing science data for downlink. Scientists evolve goals by modifying onboard software as a better scientific understanding of the target emerges throughout the mission.
- Autonomy architectures and software engineering is in many ways the glue that binds together all the capabilities listed above. This area addresses basic separation of reasoning engines from models and knowledge, the design of modeling languages and development of modeling tools, code and test generation, specific autonomy software testing concepts, and architectures and development environments that promote easy, flexible software reuse from mission to mission.

Most of these autonomy capabilities are being developed now as part of an initial emphasis on autonomy for spacecraft or engineering functions. Such capabilities directly address loop-closing and cost reduction goals. But as time goes on, autonomy development will be

² Power is an example of an always-scarce onboard resource which must be carefully validated so as not to be oversubscribed. Some spacecraft activities must happen within brief time windows to be meaningful, i.e., observing the natural satellite of a planet at closest approach point of a trajectory or orbit.

³ Spacecraft safing is highly desirable from a reliability viewpoint, but does result in the mission being suspended while the spacecraft awaits instructions from Earth. In some situations it is the *wrong* thing to do. For example, a safing response during orbit insertion results in a working spacecraft but a lost mission.

targeted more and more towards serving the science side of the missions. That work has begun even now. Once a critical mass of autonomy capability is in place, it can also be expected that there will take place an intersection with other computer science technologies. In particular, it is easy to imagine scenarios where alerts based on science event detections on remote space platforms are downlinked, then broadcast over a future, extended version of the World Wide Web where not only NASA scientists, but members of the general public can receive light-delayed, but otherwise real-time imagery of volcanic eruptions on Jupiter's moon Io, and the like. Indeed such a scenario seems almost a logical conclusion of the technology development taking place in autonomy and other areas right now.

Common Concerns and Issues

The notion of spacecraft autonomy raises concerns and issues in the minds of many people at NASA, about technological maturity, about risk, about feasibility from a systems engineering viewpoint, about actual benefits. Here, some of these concerns are enumerated, and short responses are provided. Neither the list of concerns nor the responses should be taken as complete or final. It is important to reemphasize that the emergence of spacecraft autonomy at NASA is taking place against a general background of cultural change, but the questions concerning autonomy have already moved beyond "why?" to "how?"

A common concern is an example of a systems engineering issue: Will there be adequate computing resources onboard future spacecraft to support the more sophisticated flight software that is implied by autonomy? The answer to this concern appears to be a relatively recent "yes." In parallel with autonomy technology development, NASA is also pursuing aggressive technology development in the areas of flight computers and memory. While this concern might have been a show-stopper only a few years ago, it is now anticipated that scaleable processors in the 100+ MIPS range and gigabytes of onboard storage will be routinely available for future missions. Such specifications are well within the real-time and footprint needs of autonomy software currently under development. New forms of software fault tolerance are being developed as well, to contribute to solving the problem of operating in high-radiation environments, usually approached as a hardware fault tolerance problem only.

The communications resource is more interesting. As noted above, instrument and sensor technologies routinely advance the capacity for collecting data onboard space platforms. Concurrent technology development in communications, particularly in optical communications, will help to offset this trend by increasing link bandwidth capacity. However, the situation is a perfect example of race conditions and will probably never be eliminated. Given this, it is almost certainly the case that abilities for performing onboard science data analysis to either prioritize downlink or intelligently summarize science data will also play an important role in addressing this particular resource challenge.

Another concern has to do with whether autonomy development will really lead to cost reductions. It is pointed out that first use applications of new technology rarely provide cost savings. This is true, and the straightforward response is that technology development costs have to be amortized across several mission uses before the savings is apparent. But there is a different, and more subtle answer to this concern as well.

When new concepts and technologies are introduced in one part of the mission lifecycle, often new costs appear elsewhere in the lifecycle, in a strange kind of manifestation of an apparent conservation law. The way to prevent this phenomenon is to introduce new concepts and technologies *across* the mission lifecycle, not only for their direct and complementary contributions to cost reductions, but also so that there is completeness, and

no easy cracks for new costs to fall through. Without such awareness, autonomy software testing could easily represent one of these cracks. Autonomy software, which is intended to support reasonable decision-making in scenarios which have *not* been anticipated, cannot be meaningfully tested with a scenario-enumeration or even a scenario-sampling approach. New testing concepts and approaches will be needed⁴. Fortunately, autonomy work has contributions to make to design, development, integration and test, and of course, operations. The best way to realize the cost reduction potential of autonomy is to apply new ideas in software engineering, and probably systems engineering as well, right across the mission lifecycle.

Yet another concern has to do with the perceived additional risk implied by any new technology. One response to this concern is to note that a technology does not entail risk in and of itself, but it is rather how it is used that determines the level of risk. To give a specific example, science autonomy developments suggest that the critical downlink resource might be usefully partitioned on future spacecraft between raw data, data that has matched a recognizer, and data which has passed some form of “interestingness” measure and needs to be examined by a scientist as a candidate discovery. The choice of how to weight the use of these downlink partitions is up to the mission designers and scientists, and in fact the technology may be used differently, perhaps more boldly as the mission unfolds, for several reasons: more confidence in the technology, the primary science goals for the mission have been achieved, there is a better basis for using recognizers, there is reduced support for continuing the mission, etc. The point is that the technology provides more options and flexibility, but risk posture is still for mission and science personnel to decide upon.

The risk issue for autonomy also takes the form of concern about loss of predictability of spacecraft events, or equivalently, loss of precise tracking of spacecraft state. Strictly speaking, this observation is true, but it typically ignores the reasons why it is true. Autonomy software consciously takes into account the onboard context in which activities are to be carried out; this context can include not only spacecraft internal state, but also the environment. This property of autonomy software makes it difficult to test, most certainly, but it also targeted towards an unprecedented form of robustness which traditional spacecraft sequences do not provide. Autonomy software can be resilient, continuing to try to find alternate ways of executing commands and achieving mission goals despite execution glitches, faults, and other unanticipated events. Traditional sequences may safely preserve the spacecraft, but the mission gets interrupted, pending ground intervention, when a sequence or contingency cannot execute properly. The flip side of unpredictability is effectively grappling with uncertainty, and this is much of the promise of autonomy. The autonomy technology developers at NASA fully acknowledge that this robustness property of autonomy software has not yet been convincingly demonstrated. However, the future *in situ* missions all involve space platforms interacting directly with their environments, raising the stakes on the amount of uncertainty that will have to be dealt with. Autonomy does imply a trade between predictability and robustness in execution, but it is a well-considered trade, and an appropriate one for the times, in light of the nature of the future missions.

Organization and Scope of this Issue

This special issue on Autonomous Space Vehicles will focus on the work ongoing at NASA on developing autonomy technology for NASA’s spacecraft missions. Autonomy technology is being developed in other contexts as well, notably for mobile robot applications. It is also the case that NASA is not the only government agency with interest

⁴ See the article by Lowry and Dvorak in this issue.

in autonomous space vehicles. Research and technology development on autonomy is being conducted out of other government agencies, many university laboratories, and industry as well [Doyle et al 98]. However, the scope of the articles presented here cover work being done to address NASA's unique set of drivers for achieving autonomy on space platforms to realize its future mission set.

Within NASA, autonomy technology development is mostly centered at Ames Research Center (ARC) and the Jet Propulsion Laboratory (JPL), the two NASA Centers where critical mass has long existed in AI research [Williams and Nayak 96a, Chien et al 97]. Autonomy efforts are also taking shape at Goddard Space Flight Center (GSFC) and Johnson Space Center (JSC).

The largest effort in spacecraft autonomy development at NASA currently is the Remote Agent, a joint technology project by ARC and JPL [Bernard et al 98, Pell et al 98]. The Remote Agent Experiment will be conducted on the New Millennium Deep Space One spacecraft in late 1998, a mission being readied now at JPL for a July 1998 launch, whose primary goal is to flight validate new technologies.

The Remote Agent consists of a Smart Executive [Pell et al 97], a Planning and Scheduling module [Muscettola et al 97], and a Mode Identification and Reconfiguration (MIR) module [Williams and Nayak 96b]. The onboard system receives mission goals as input, which are translated to a set of spacecraft activities free of resource and constraint violations by the Planner/Scheduler. The Smart Executive provides robust, event-driven execution and runtime monitoring and decision making. MIR continuously monitors qualitative representations of sensor data, identifying current spacecraft modes or states, and when these are fault modes, selects recovery actions. Other functions such as guidance, navigation and control, power management, and science data processing are domain-specific functions that can be layered on top of this basic autonomy architecture, and are developed or modified for each new mission. The Remote Agent has been designed to be a core architecture for autonomous spacecraft.

This issue features articles on component technologies of the Remote Agent (Chien et al, Gat and Pell) which report on the specific form of the technology used in the Remote Agent, its AI research heritage, and other applications of the technology within NASA.

In addition, the article by Gamble and Simmons looks at the Remote Agent as a case study in the space of possible autonomy software architectures. The challenge here is to balance software engineering goals, particularly reuse considerations, against a wide range of specific NASA mission needs for autonomy.

The Remote Agent directly targets autonomy for engineering functions of the spacecraft spanning mission planning, resource management and fault protection. As noted above, onboard autonomy to process and analyze science data will be equally important to NASA's future missions. This work has begun, is based on image analysis, machine learning, knowledge discovery and data mining techniques [Cheeseman et al 96, Stolorz and Dean 96], and is reported on in the article by Stolorz and Cheeseman.

One of the most vital issues concerning autonomy has to do with how to test and validate autonomy software. This is a central challenge, raised beyond the normal challenge of validating flight software by the fact that autonomy software is meant to make closed-loop decisions in uncertain contexts. The article by Lowry and Dvorak speaks to this important area, describing approaches based on formal methods, AI techniques, and software engineering common sense.

Even given the NASA focus of this special issue, the survey of autonomy work at NASA to be found here will be incomplete. Some of the other projects in NASA autonomy technology development are described in [Doyle et al 98].

The Missions of Exploration

As has been argued here, autonomy has strategic importance for many of the missions NASA has planned for the future. NASA missions are organized into three so-called Enterprises: Space Science, with primary responsibility at JPL; Earth Science, with primary responsibility at GSFC; and Human Exploration and Development of Space (HEDS), with primary responsibility at JSC. The three mission sets impose different kinds of drivers on autonomy technology development. In the Space Science mission set, the central difficulties associated with light-time delayed and tenuous communication, coupled with the sparse prior information available on deep-space planetary targets make the need for autonomy to respond, in context, to unanticipated engineering and science events fairly obvious and imperative. This is particularly the case in the upcoming wave of *in situ* missions where direct interaction with a remote planetary environment adds more uncertainty to what is already largely unknown. Planetary exploration (and someday, extra-solar system exploration) will always place the most severe demands on autonomy. For this reason, the majority of the mission examples given here are drawn from the Space Science Enterprise, which in no way diminishes the importance of the contributions to be made by autonomy to the Earth Science and HEDS Enterprises, or of the effort being placed there.

The looming challenge in the Earth Science Enterprise is grappling with truly overwhelming amounts of data -- on the order of terabytes a day -- which will be collected and downlinked from fleets of Earth observing space platforms. Another challenge is automated planetary monitoring for hazards such as forest fires, volcanic eruptions, and poorly understood phenomena like El Niño. See Figure 1. Earth orbit is also the first place where formations and constellations of spacecraft will appear -- with their attendant control and coordination challenges.

The driving consideration in the HEDS Enterprise is to find the right ways to combine human and machine intelligence into a single, effective system. One of the unique challenges is to automatically track state accurately enough when a human enters a control loop so that the updated context can be made available once control reverts to the machine -- a kind of cognitive clutch. Any applications of autonomy in the HEDS Enterprise will be always stringently evaluated against human safety concerns.

We now turn in this final section to a quick survey of some of the fascinating upcoming missions, describing their science and exploration goals -- many of which are unprecedented -- and examining specifically what autonomy has to offer.

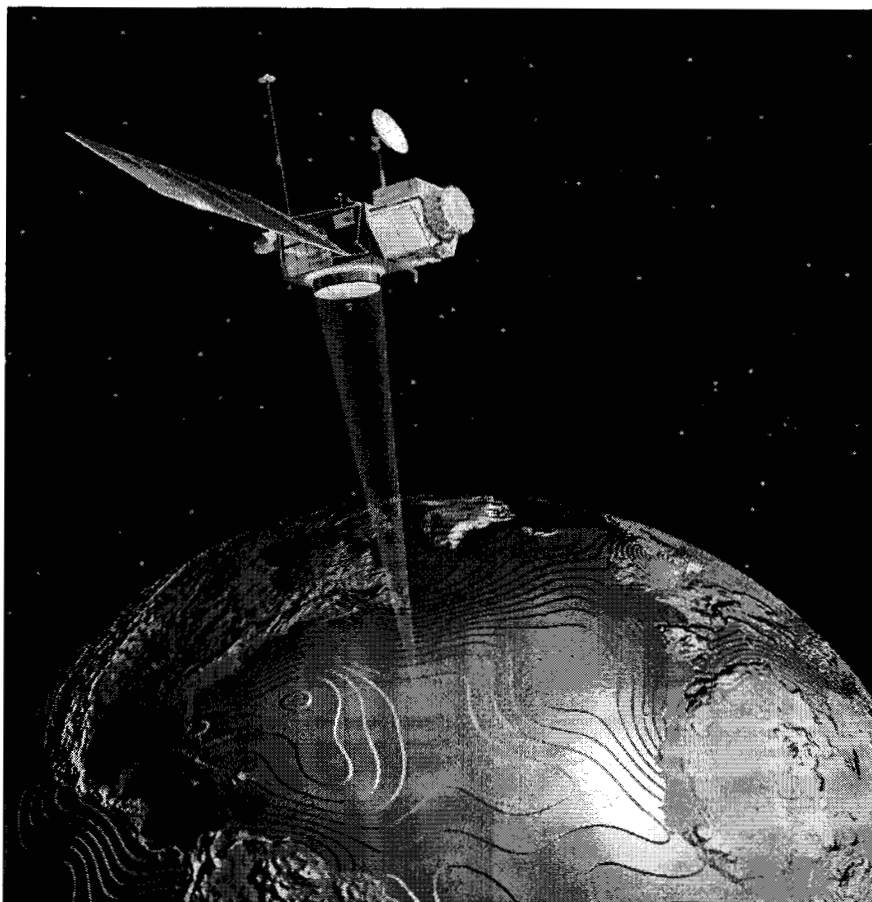


Figure 1. Automated Analysis of Earth Observing Data.

The Mars 2003 and Mars 2005 missions will both return rovers to the surface of Mars. These missions will have more ambitious goals than Pathfinder / Sojourner: in the number of sites to be investigated, the breadth and depth of science investigations to be conducted, and the total amount of terrain to be traversed. See Figure 2. The basic mission goal in each case is to collect and cache a sample of Mars rocks and other surface material (one or the other cache will be retrieved and returned to Earth as part of the Mars 2005 mission), performing *in situ* analysis both to support the selection of cache material and to return intermediate data in the normal way during the missions. The '03 and '05 rovers will each carry a full complement of scientific instruments and sensors to, among other goals, continue the investigation of conditions and possibilities for life on ancient Mars.

The rovers will operate in two major modes: conduct science investigations at a site, and traverse between sites. One important use of autonomy to maximize the scientific return of these missions is to have the onboard capability to interrupt a traverse based on the detection of scientifically interesting phenomena (outcroppings, unusual mineralogical signatures, evidence of water). The rover should keep its head up while moving! Another important use of autonomy is to adapt the performance of the rover by learning models of rover performance in the Martian environment. Even a few percent increase in locomotion efficiency and resource usage can translate into significant additional scientific throughput when integrated over the entire mission.

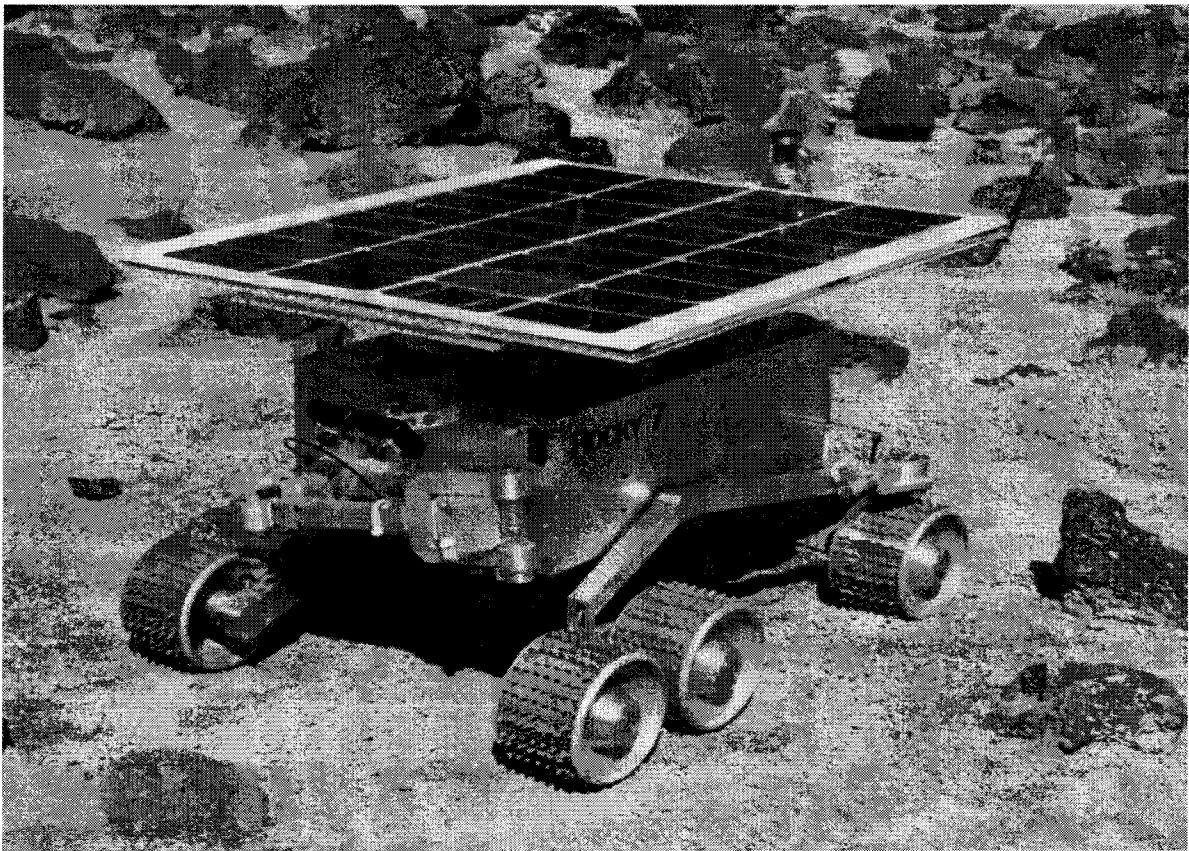


Figure 2. Long-traverse Mars Rover Missions.

The Europa Orbiter mission, which is slated for launch in 2003, will perform focused investigations of this most intriguing of Jupiter's moons. Europa fires the imagination because of current theories on the existence of a subsurface ocean. Tidal effects due to the proximity of immense Jupiter and orbital resonances among the Jovian satellites exert forces of considerable magnitude at Europa, great enough perhaps to release the thermal energy which could result in a layer of liquid water beneath the surface (Europa has long been known to be mostly a water-ice object, from Earth-based spectroscopic studies). Recently, organic material has been detected on the surface of Ganymede and Callisto, two of the other Jovian satellites, raising the stakes further on the possibilities for Europa to harbor the three basic ingredients of life: water, an energy source, and organic material.

Europa has a dramatically disrupted surface, and one of the forms of indirect evidence for the existence of the subsurface ocean is the scale of tectonic movements on the European surface. See Figure 3. Autonomy can help here. The Europa Orbiter spacecraft can arrive with archived image data of the surface of Europa from the previous Voyager and Galileo missions. The spacecraft will also begin to collect new data which can also be archived onboard. Then there is a local basis -- at three different time scales -- to detect change on the surface of Europa. If such evidence of tectonics is found, the specific images can be tagged for high-priority downlink, in a natural and compelling example of using onboard data analysis to pursue science goals while efficiently addressing the constraints of deep-space communications.

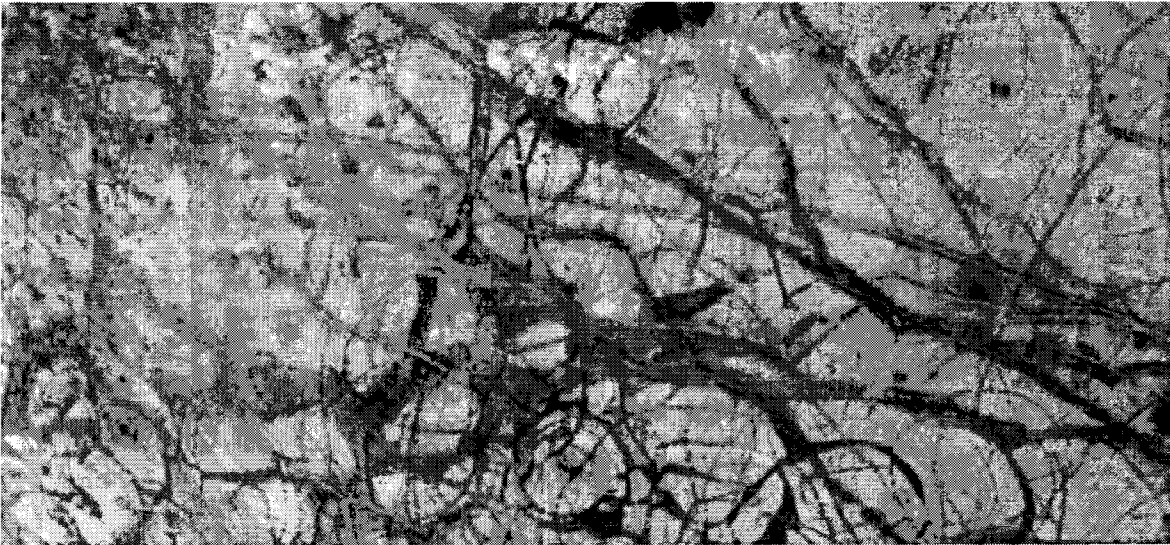


Figure 3. Change Detection on Planetary Surfaces.

The Origins program is a new NASA program whose goals are to investigate the ultimate origins -- of the universe, of galaxies, of life. The Planetfinder mission within this program may turn out to be its flagship mission. Planetfinder will be a deep-space interferometer most likely comprised of several elements. Using interferometry to null the light coming from nearby stars (out to 50 light-years), and then systematically search for planetary companions of those stars, this mission has the goal of directly imaging Earth-class planets by 2010 or so and ultimately resolving continental masses on their surfaces. The search for life in the universe has recently taken a number of palpable and exciting forms at NASA.

The need for autonomy on Planetfinder stems from the multiple-platform aspect of the mission. The interferometer would be composed from several spacecraft elements and a special challenge results from the need to perform pointing of the entire formation with unprecedented precision for truly deep-space observing. See Figure 4. If this collective platform is to be operated at low cost, then the inevitable -- and divergent -- degradations of performance which will appear over time across the distinct platforms must be automatically detected, evaluated, and compensated for to preserve the overall coordinated pointing accuracy of the interferometer. On the science side of this mission, automated classification of detected planets is a possibility, as is automated spectroscopic analysis of atmospheric constituents of Earth-like planets.

Before the Planetfinder mission is realized, formations and constellations of spacecraft in Earth orbit will appear, with objectives for earth observing (natural event detection, atmospheric and oceanographic studies, land-use and ecological management), and communications (networks such as Iridium and Teledesic). The salient difference between formations and constellations is whether the individual satellite assets are similar or not, and whether a strict geometric configuration is required to perform the mission. In general, homogeneous formations are appropriate to support low-earth orbit (LEO) satellite-based communication networks, while heterogeneous constellations provide more and desired flexibility for Earth-observing objectives.

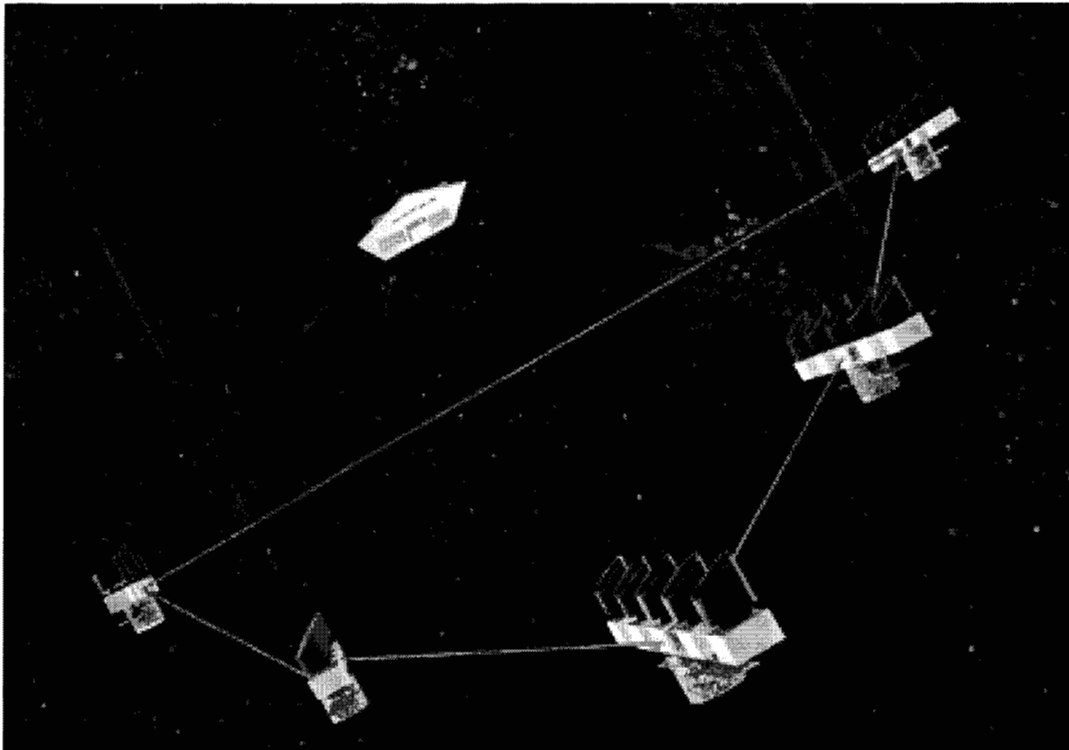


Figure 4. Deep-sky Interferometry.

In formations and constellations, spacecraft functions become distributed, and do not simply scale from the single-platform case: this applies to mission planning, resource management, execution, and fault protection, as well as to information sharing and problem solving. Multiple-platform missions also require shared approaches to operations to keep costs down, with, for example, ongoing engineering data summarization and paging alerts when problems occur. Finally, automated orbit maintenance, including onboard navigation, maneuver planning and execution, along with de-orbiting of compromised satellites and automated promotion of satellites held in reserve, will be needed to maintain formations at low cost.

One area of the manned space program where autonomy concepts are being looked at in earnest is in the next-generation Space Shuttle concept. The challenge here is to reduce both the cost and the turnaround time associated with the flight of the vehicle. The specific goal is to slash payload costs by a factor of ten and to achieve a routine seven-day turnaround. A number of different designs have been examined recently, with the one known as X-33 going forward to detailed design and initial flight tests.

Autonomy figures prominently in the emerging NASA concept to operate a low-cost, quick-turnaround reusable launch vehicle for low Earth-orbit manned missions: Onboard software conducts ongoing fault and performance monitoring. Salient engineering data is downlinked automatically and requests for maintenance and repair are also generated automatically. Such requests are input to a ground-based automating planning and scheduling system which generates and updates a maintenance plan and schedule for refurbishing the vehicle even while it is in flight, for immediate execution upon landing.

Returning to deep space, a mission which may fly by the year 2004 is the Pluto/Kuiper Express mission. Pluto is the only known planet which has yet to be visited by a

spacecraft. Historically, Pluto has been something of an enigma. The first four planets (including Earth) are small and rocky, with thin atmospheres. The next four, known as the gas giants, are large and may be almost entirely composed of gas. Pluto, despite its great distance, seemed more like the terrestrial planets than the gas giants. This mystery may now be solved, with the emerging understanding of a third class of objects in the solar system, the so-called Kuiper objects, of which Pluto may be the most outstanding member. A mission to Pluto is now even more compelling in the context of this new theory.

Any trajectory to Pluto is dominated by the extremely long cruise period required to reach the most distant planet. The Pluto Express mission calls for on the order of a twelve-year cruise, and this includes the benefit of a gravity assist at Jupiter. To keep costs reasonable, Pluto mission personnel conceived an innovative operations concept known as Beacon Operations. On a continuous basis, the spacecraft sends a simple signal which denotes the urgency with which interaction with the ground is needed. This concept assumes a certain level of autonomy on the spacecraft, certainly for fault protection, but perhaps for detecting science events as well. The Beacon Mode Operations concept includes the idea of onboard engineering data summarization in an ongoing fashion, so that when an emergency signal is received from the spacecraft, it is quickly followed -- once a full communications link is established -- by an anomaly report, including context and completed analysis, to bootstrap the ground-based troubleshooting effort.

Perhaps the mission currently on the NASA books which cries out for autonomy more than any other is the Deep Space Four (DS-4) mission, which is the rather generic name given to a mission with a planned 2002 launch which is to rendezvous with, land on, and return a sample from a comet. See Figure 5. Comets are scientifically intriguing in that they are thought to contain primordial material largely unaltered from the era of the formation of the solar system. DS-4 would rendezvous with its comet at the range from the sun where interaction with the solar wind begins to produce noticeable activity -- the beginnings of the tail.



Figure 5. Sample Return from a Comet.

What makes the DS-4 mission so intriguing from an autonomy viewpoint is the extreme unpredictability of the cometary environment. Comets can spontaneously emit jets, eject particles, even break up. A mission to rendezvous with, much less land on, a comet must be able to detect events which represent potential hazards to the spacecraft and mission, as well as being science events in their own right. The set of onboard autonomy capabilities which appear relevant here, at a minimum, are event and hazard detection, object tracking, navigation, and maneuver planning and execution. These capabilities must be tightly integrated so that decision loops can close quickly, to, for example, abort landings or execute safety maneuvers.

Aerobots are a newly conceived space platform concept which combines the wide coverage advantages of orbiting spacecraft with the *in situ* exploration advantages of surface vehicles such as rovers. The basic idea is to exploit the diurnal thermal cycle of a planetary environment to alternately go aloft into the prevailing winds and land on the surface (where there is one), once a day (or sol, the equivalent in the local planetary environment). A planetary hot air balloon with a serious scientific payload. The concept works wherever a planetary atmosphere exists, including at Venus, Mars, Jupiter, and Saturn's moon Titan.

An aerobot would require a high degree of onboard autonomy, because entering a dense atmosphere (Mars being the exception among the examples given above) implies difficult communications, with much of the mission being performed without routine interaction with the ground. Aerobots also suggest a unique form of the path planning problem: presumably the vertical dimension of motion can be reasonably controlled, but the two horizontal dimensions will have a significant stochastic element, and path planning would have to be based on models of wind patterns. This suggests an approach of arriving with crude wind models derived from Earth-based observations, and refining them based on actual experience in the planetary atmosphere. In aerobot missions, scientists might experience a certain frustration with analysis results not achieved *in situ*, because it would be nearly impossible to return to a site.

Perhaps the mission with the most remarkable set of stretch goals -- for both autonomy and general engineering functions -- is the proposed Europa cryobot / hydrobot mission. This mission would land on the surface of Europa, melt through its icy crust, and release an underwater submersible into the suspected subsurface ocean. See Figure 6. The problems to be solved are mind-boggling. First of all, melting through perhaps several kilometers of ice at a temperature which gives ice the structural properties of rock -- and starting from vacuum -- is unprecedented. Going tethered or untethered each present unique challenges. A tethered mission would solve the communications problem, but reaching the European ocean floor, which may be a hundred kilometers from the ice/water boundary, becomes problematic. On the other hand, going untethered forces one to look at acoustic communication within the ocean, or navigating back to the penetration site, or somehow reemerging through the ice crust at a different site.

The need for autonomy on this mission is obvious, for the usual drivers of poor communications and an uncertain environment are multiplied many times. It's hard to imagine sending a spacecraft into a more alien environment. And yet, it's also hard to imagine a more compelling place to explore. Nowhere else in our solar system have we any reason to expect to find an ocean, perhaps the defining global characteristic of our own planet. We've already noted that Europa may harbor the basic ingredients of life. Would we be able to equip our intelligent envoy to know what to look for, and the means to recognize it?



Figure 6. Exploration of Unknowable Environments.

The above quick survey is just a sample from the incredibly exciting set of future NASA missions. The space agency is experiencing a return to its most noble goals of exploration: the search for life in the universe, and a new vision of sustained, vigilant intelligent presence in the solar system and eventually beyond, via a fleet of autonomous space vehicles. Autonomy done well means tapping the expertise not only of computer scientists, but of spacecraft engineers, mission designers, operations personnel, software engineers and systems engineers. But for the first time, AI practitioners will work side by side with these traditional contributors, to realize the future NASA mission set.

Acknowledgments

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[Bernard et al 98] Bernard, D, G. Dorais, Chuck Fry, E. Gamble Jr., R. Kanefsky, J. Kurien, W. Millar, N. Muscettola, P. Nayak, B. Pell, K. Rajan, N. Rouquette, B. Smith, and B. Williams, "Design of the Remote Agent Experiment for Spacecraft Autonomy", in *Proc. of IEEE Aeronautics Conference (Aero-98)*, Aspen, CO, IEEE Press, 1998.

[Cheeseman et al 96] Cheeseman, P, B. Kanefsky, R. Kraft, J. Stutz and R. Hanson, "Super-resolved surface reconstruction from multiple images", in *Maximum Entropy and Bayesian Methods*, G. R. Heidbreder ed., Kluwer Academic (1996).

[Chien et al 97] Chien, S., D. DeCoste, R. Doyle and P. Stolorz, "Making an Impact: Artificial Intelligence at the Jet Propulsion Laboratory," *AI Magazine*, Spring 1997.

[Doyle 97] Doyle, R. "The Emergence of Spacecraft Autonomy," Invited Talk, *14th National Conference on Artificial Intelligence*, Providence, RI, July 1997.

[Doyle et al 98] Doyle, R., R. Rasmussen, G. Man, and K. Patel, "Report on the Highly Autonomous Systems Workshop," *AI Magazine*, to appear.

[Muscettola et al 97] Muscettola, N., B. Smith, S. Chien, C. Fry, K. Rajan, S. Mohan, G. Rabideau and D. Yan, "On-board Planning for the New Millennium Deep Space One Spacecraft," *Proceedings of the 1997 IEEE Aerospace Conference*, Aspen, CO, February, 1997.

[Muscettola et al 98] Muscettola, N., P. Nayak, B. Pell, and B. Williams, "Remote Agent: To Boldly Go Where No AI System Has Gone Before," *Artificial Intelligence*, 103(1/2), August 1998.

[Pell et al 97] Pell, B., E. Gat, R. Keesing, N. Muscettola and B. Smith, Plan Execution for Autonomous Spacecraft, *15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan, August 1997.

[Pell et al 98] Pell, B., D. Bernard, S. Chien, E. Gat, N. Muscettola, P. Nayak, M. Wagner, and B. Williams, "An Autonomous Spacecraft Agent Prototype," *Autonomous Robots*, to appear.

[Stolorz and Dean 96] Stolorz, P. and C. Dean, "QuakeFinder - A Scalable Datamining System for Detecting Earthquakes from Space", *2nd International Conference on Knowledge Discovery and Datamining*, Portland, OR, August 1996.

[Williams and Nayak 96a] Williams B., and P. Nayak, "Immobile Robots: AI in the New Millennium," *AI Magazine*, Fall 1996.

[Williams and Nayak 96b] Williams, B., and P. Nayak, "A Model-based Approach to Reactive Self-Configuring Systems," *13th National Conference on Artificial Intelligence*, Portland, OR, August 1996.